Shape From Motion of Nonrigid Objects: The Case of Isometrically Deformable Flat Surfaces

Ricardo Ferreira
ricardo@isr.ist.utl.pt

João Xavier
jxavier@isr.ist.utl.pt

João Costeira
jpc@isr.ist.utl.pt

Instituto de Sistemas e Robótica
Instituto Superior Técnico
Lisbon, Portugal

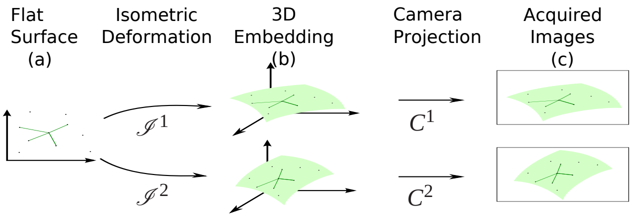

Figure 1: Acquisition model of the isometrically embedded surface observed by a camera.

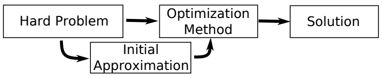

Figure 2: Overview of the optimization strategy. First an initial approximated solution is computed to start the optimization.

The purpose of this paper is to allow reconstruction of deformable surfaces isometrically embedded in 3D, e.g. a flag waving at the wind or someone waving a sheet of paper, from image data. It is assumed that a set of non-calibrated images is available with matched features between them. For simplicity the cameras here are restricted to be scale-orthographic, each modeled by a single scale factor and extrinsic parameters. Figure 1 illustrates the acquisition process. From a single flat surface, modeled as a set of features, 3D isometric embeddings are generated by passing the feature points through embedding functions $\mathscr{I}^k$. These are then viewed by different cameras yielding the observed images.

There are two problems that can be formulated with images of isometrically embedded surfaces. The first consists of estimating the 3D embedding (pose) of the surface in a particular image k, when the flat surface is assumed to be known. In figure 1 this roughly means estimating (b) given (a) and (c). This problem will be called the pose estimation problem, but will not be the focus of this paper. The second problem is given several images estimate the surface that generates them. In the figure this means from several observations (c), obtain (a). Here this problem shall be named the surface estimation problem and the embeddings (b) are not considered important. Although not done yet, a future objective will be to unify both problems, i.e. estimate the generating surface (a) and the various embeddings (b) given only the set of images (c). This can be trivially achieved by first estimating the surface and then applying a known pose estimation algorithm, future work will focus on integrating both in a single problem.

This paper models accurately flexible flat surfaces following a similar strategy presented in [1] where an initial approximation is given by a sequential algorithm. This approximation is then fed to a global cost function, further refining the result and hopefully converging to the global minimum. To better capture global constraints instead of just measuring local fit in which integration error can accumulate a different optimization function is used. Here we deal with the more realistic assumption that data can be missing, which is of utmost importance in deformable surfaces where self occlusion and partial observations are common. This paper handles missing data and provides the required performance tests. Finally, one important contribution is how to obtain second order information about the 3D embeddings, such as how much the surface bends, which can be used for measurements and reused in the cost function. Although the surfaces are assumed to be locally planar, second order information provides relevant information in the presence of a sparser data set, as so often occurs.

In a real world example, 12 images of a bed cover were taken (see figure 3) at various angles and differently folded. In these images, 118 different points were hand clicked (when visible) and the algorithm was run on them. The results obtained are shown in figure 4 (overlaid on a picture taken of the flat cloth fabric). Results provide a benchmark for real world data, hand clicked, not very dense for the amount of bending, and for embeddings not truly obeying the isometric properties since cloth is easily sheared.

Figure 3: Images (3 out of 12) taken of a bed cover made of cloth.

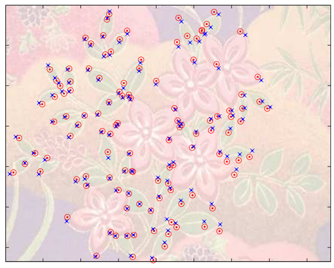

Figure 4: Reconstruction of the bed cover cloth overlaid on an image taken of the cloth laying flat. Blue crosses are the reference clicked points, red circles are the results given by the algorithm.

When a flat surface is embedded in 3-D one important property is verified: through every point of the embedded surface there is a direction where it is locally linear. This means that through each point, the embedding is allowed to curve in only one direction, making it locally like a cylinder (this is a second order approximation). Since the word “curvature” has a very precise meaning (all the surfaces here considered have 0 curvature), here this embedding-specific second order property shall be called “bend” at a point, and its “bend radius” will be the radius of the smallest osculating circle through that point. In differential geometry this terminology refers to the sectional curvature.

Bend radius and axis were computed for an embedding wrapped around a cylinder (bend radius at each point is 1). Figure 5 shows the obtained results. Note that although at most points the bend axis is similar, there are a few outliers. The histogram shows that bend radius estimates concentrate around the correct value but some noise exist.

[1] Ricardo Ferreira, João Xavier, and João Costeira. Reconstruction of isometrically deformable flat surfaces in 3d from multiple camera images. *ICASSP*, 2009.

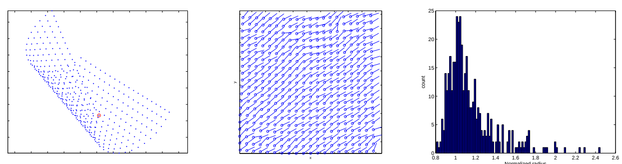

Figure 5: Results of estimating the bending axis and radius from the data. Left: Image to be measured. Middle: The computed axis of bend at each point is shown. Right: histogram of the computed radius at each point (the radius is normalized to be 1 unit).